Reducing the effect of seismic noise in LIGO searches by targeted veto generation

D. M. Macleod[1], S. Fairhurst[1], B. Hughey[2], A. P. Lundgren[3,4], L. Pekowsky[4], J. Rollins[5], J. R. Smith[4,6]

[1]Cardiff University, Cardiff, CF24 3AA, United Kingdom
[2]University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA
[3]The Pennsylvania State University, University Park, PA 16802, USA
[4]Syracuse University, Syracuse, NY 13244, USA
[5]LIGO - California Institute of Technology, Pasadena, CA 91125, USA
[6]California State University Fullerton, Fullerton CA 92831, USA

E-mail: duncan.macleod@astro.cf.ac.uk

Abstract. The Laser Interferometer Gravitational-Wave Observatory forms part of the international effort to detect and study gravitational waves of astrophysical origin. One of the major obstacles for this project with the first generation detectors was the effect of seismic noise on instrument sensitivity – environmental disturbances causing motion of the interferometer optics, coupling as noise in the gravitational wave data output. Typically transient noise events have been identified by finding coincidence between noise in an auxiliary data signal (with negligible sensitivity to gravitational waves) and noise in the gravitational wave data, but attempts to include seismometer readings in this scheme have proven ineffective. We present a new method of generating a list of times of high seismic noise by tuning a gravitational wave burst detection pipeline to the low frequency signature of these events. This method has proven very effective at removing transients of seismic origin from the gravitational wave (GW) data with only a small loss of analysable time. We also present an outline for extending this method to other noise sources.

1. Introduction

The Laser Interferometer Gravitational-Wave Observatory (LIGO) [1] is designed to detect and study gravitational waves (GWs) of astrophysical origin. The project operates three detectors at two sites in the United States: the LIGO Hanford Observatory (LHO) in Hanford, Washington, hosting both the H1 and H2 detectors, and the LIGO Livingston Observatory (LLO) in Livingston, Louisiana, housing the L1 detector. Each of these is a kilometer-scale, power-recycled, Michelson laser interferometer with Fabry-Perot arm cavities. Together with the GEO600 [2] and Virgo [3] detectors, they form a global network of GW interferometers operated in a collaborative effort towards the detection of GWs.

The three Initial LIGO (iLIGO) detectors reached their design sensitivity during the fifth science run (S5), which ran from November 2005 until November 2007, before upgrades were installed as part of the Enhanced LIGO (eLIGO) project [4]. H2 was not upgraded and did not take part in the sixth science run (S6), which ran from July 2009 until October 2010. These machines are involed in the search for GW signatures from, in particular, the coalescence of compact binary systems containing black holes or neutron stars [5], and unmodelled burst events [6], for example supernovae. There are also a variety of searches for longer-duration signals that are not discussed in this work.

The output of each LIGO detector is a single data stream, $s(t)$, that in general contains some combination of a GW signal, $h(t)$, and detector noise, $n(t)$,

$$s(t) = h(t) + n(t). \tag{1}$$

Various kinds of transient noise events (glitches) with instrumental or environmental origin can mask or mimic true astrophysical signals, thus limiting the sensitivity of any search that can be performed over these data [7, 8, 9].

In the searches for short-duration GW signals the noise background is dominated by glitches that are difficult to separate from short-lived astrophysical signals, and although there are techniques to distinguish between the two, for example multi-detector coincidence or signal consistency tests, it is preferable to understand the physical origins and eliminate them. Throughout the lifetime of LIGO up to and including S6, search sensitivity has been improved by careful use of *vetoes*: time segments indicating poor data quality (DQ) which are removed from the analysis. This strategy has allowed data analysts to tune and operate search pipelines using a subset of cleaner data, thus increasing the chance of extracting a signal from the noise [10].

The detrimental effect of seismic noise has been known to be a key limiting factor to the sensitivity of GW detectors at low frequencies (below $\sim 70\,\mathrm{Hz}$ [11]). However it is also a common cause of glitches due to non-linear coupling of low-frequency seismic noise to higher-frequency noise in the gravitational wave readout. Previous methods to generate veto segments for times of high seismic noise have proven ineffective.

In this paper, we introduce a new method of constructing vetoes for the specific case of seismic noise that has proven highly effective when used in the latest searches

Frequency (Hz)	Distance (km)	Source
$0.01 - 1$	10^3	Distant earthquakes Microseism
$1 - 3$	10^1	Far anthropogenic noise Close earthquakes Wind
$3 - 10$	10^0	Anthropogenic noise Wind
$10 - 30$	10^{-1}	Close anthropogenic noise

Table 1. Description of the main seismic frequency bands and their sources

for transient GW signals. As will be shown, this method can be extended to the study of other environmental and instrumental systems in order to maximise the performance of vetoes for the Advanced detector era.

The paper is set out as follows. In section 2 we will outline the seismic environment at each of the LIGO sites, and the effect it has on detector sensitivity. In section 3 we describe the existing methods used in S5 and S6 and their shortcomings for the application to seismic noise. In section 4 we describe our procedure for generating lists of events in seismometer data and constructing veto segments around them. In section 5 we present the results in terms of veto efficiency and deadtime. Finally, section 6 presents a brief discussion of further applications of the method.

2. Seismic noise in LIGO

In this section we will briefly characterise the seismic environment at each site, and its effect on the LIGO interferometers. The various types of seismic noise, as characterised by their source, can be separated into four frequency bands as given in Table 1. Their contribution to the overall LIGO strain sensitivity curve (known as the *noise budget*) is given in Figure 1.

2.1. LIGO seismic environment

The two LIGO sites were chosen to be far from urbanised areas, thus reducing the incident seismic noise, whilst their distance from each other provides a long baseline helpful in sky-localisation of astrophysical signals [12]. They are subject to greatly varying seismic environments, each producing their own challenges for isolation.

LHO is located 15 km from the United States Department of Energy (USDOE) Hanford Site, in which several working areas include use of heavy earth-moving machinery. In addition, the Tri-Cities area begins roughly 20 km away. Both of these, along with the general use of nearby roads by vehicle traffic, contribute heavily to the

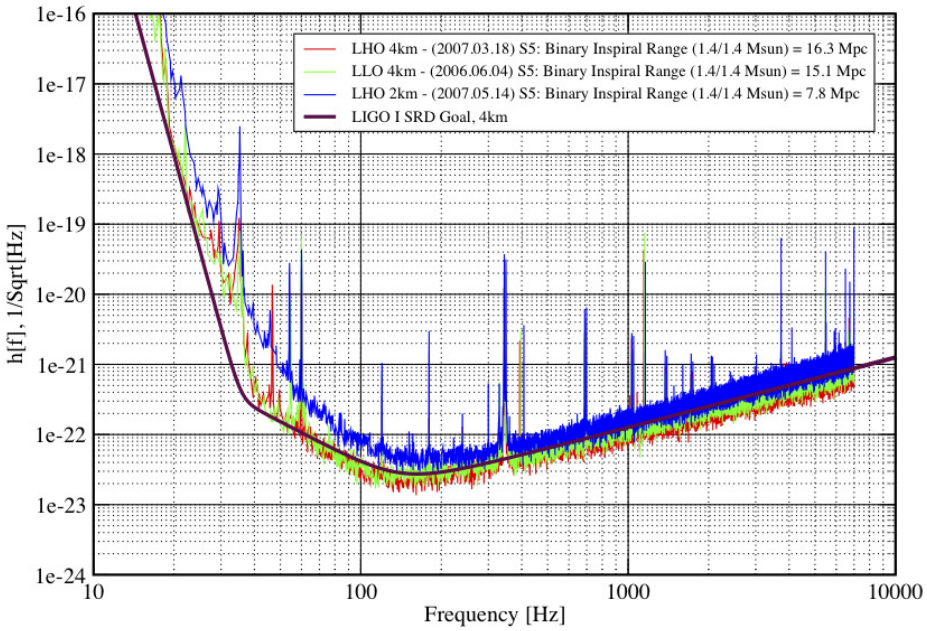

Figure 1. Strain sensitivity of the iLIGO interferometers [4]. The seismic background is the limiting noise source below ~ 50 Hz.

amount of anthropogenic seismic noise incident on the detector. The site is also subject to high winds up to 40 m/s, causing motion of the buildings and the concrete slabs supporting the instruments. These relatively local sources generate noise in the higher bands in Table 1, above ~ 1 Hz.

LLO is located 7 km from the town of Livingston, and only 3 km from a railway line used daily by cargo trains [11]. Additionally, the land surrounding LLO is used for timber harvesting. The site is only 130 km from the Gulf of Mexico, and is subject to violent rain and windstorms.

Both sites are also subject to noise from earthquakes occurring almost anywhere on Earth, and to microseismic noise from oceanic activity due to their relatively short separation from the nearest coastline. These distant events are the source of noise below ~ 1 Hz.

Due to the softer composition of the surrounding geological landscape, seismic noise was worse at LLO, so the decision was taken to install an active seismic isolation system on L1 before the fourth science run (S4). The hydraulic external pre-isolator (HEPI) feed-forward system damps low-frequency noise by using signals from the onsite seismometers to control movement of the vacuum chambers for the end test masses. This particular system was not installed at LHO before S6 – although other isolation systems were used – but will be installed as part of the Advanced LIGO (aLIGO) project [13].

2.2. Upconversion

The LIGO instruments are designed to be sensitive in the range $40 - 7000\,\mathrm{Hz}$ [1], so one may be forgiven for assuming that seismic noise below 30 Hz should not affect sensitivity in the detection band. However, *upconversion* of the low-frequency seismic motion into broadband noise has been a problem during iLIGO and eLIGO. This non-linear coupling is suspected to be in part a result of the Barkhausen effect in ferromagnetic components used for interferometer control signals. This magnetic effect generates noise in the detection band of the LIGO detectors.

3. Existing veto methods

The GW data stream is not the only information drawn from the LIGO detectors. Thousands of auxiliary data channels are recorded, containing control and error signals from instrumental systems, and measurements from the physical environment monitors (PEMs). These data are analysed in order to improve detector performance in the future, but also to identify glitches that can mimic GWs. Time segments cointaining excess noise in these channels can be vetoed from any analyses.

3.1. Veto construction

In S5 and S6, veto segments were constructed by two methods [7, 9, 10, 14].

The first method relied on known physical couplings between an auxiliary subsystem and the GW data. In this case, where a transfer function is known, the time-series output of a particular auxiliary channel is studied, and times for which a certain threshold was exceeded are recorded. Simple, but highly effective examples include overflows in analog-to-digital converters, and 'light dips' – drops in the power stored in the detector's Fabry-Perot arm cavities.

The second method replaces known couplings with statistics, applying the Kleine-Welle (KW) wavelet-based algorithm [15] to data in auxiliary channels, producing lists of events with signal-to-noise ratio (SNR) above a given minimal threshold. By design these *triggers* mimic the signal from a GW burst, and thus pollute these analyses, but come from channels with negligible sensitivity to GWs. This is verified by actuating one test mass to simulate an astrophysical signal (known as a *hardware injection*), and verifiying no signal was seen in these auxiliary channels at these times. The triggers are then tested for time-coincidence with equivalent triggers in the GW data, with a match indicating the event was unlikely to be of astrophysical origin. Veto times are chosen by selecting segments around loud triggers in the auxiliary data in order to maximise *efficiency* – the fraction of GW data triggers vetoed – whilst minimising *deadtime* – the fractional duration of data removed from the analysis. Highly effective vetoes were those with a high ratio of efficiency to deadtime.

Different implementations of this method [14, 16] were used in the searches for unmodelled bursts [6] and compact binary coalescences (CBCs) [5] during S5, producing

comparable results [7].

3.2. Shortcomings

When using data from environmental monitors, the methods above have been less effective than for instrumental signals. Time-series thresholds on auxiliary channels are very simple to construct, and can do well to identify excess noise in a given auxiliary channel. Huge efforts are put into isolating the GW readout photodetector from noise events of all kinds, and so care must be taken to measure the effects these events have on the readout, and use them to give feedback to instrumental scientists to improve hardware performance, rather than simply relying on them for vetoes.

This issue becomes especially relevant for the transition between the eLIGO instruments and those being upgraded for aLIGO [17]. It is probable that many glitch mechanisms that were being caught with time-series vetoes will be mitigated by improved hardware, thus requiring careful retuning of all time-series veto thresholds.

Furthermore, environmental systems are subject to huge variances based on the season, the day of the week, and even the time of day. Seismic noise is five times higher above $\sim 1\,\mathrm{Hz}$ during the average weekday than it is in the evening, while the increase is less than a factor of two on the weekends. It has also been seen that the LHO site is affected by the spring melt, with higher flow over nearby dams causing increased noise in the $1 - 3\,\mathrm{Hz}$ band for a period of six weeks or so each year. It would take a very complicated system of thresholds based on each of these criteria to identify noise spikes in both loud and quiet periods, which can have the same relative effect on instrument sensitivity. An example of the difference between day and night for weekdays and weekends is shown in Figure 2.

Testing for coincidence between auxiliary and readout triggers has proven very effective for instrumental channels, but, like the thresholding methods described above, not so for environmental channels. The method is centred on the KW algorithm, which was designed as a GW burst detection pipeline. As a result, this code is easily adaptable to instrumental effects where the typical frequency range mirrors that of the sensitive band of the detectors, but for seismic effects, where the typical frequencies are two orders of magnitude lower, the performance is poor. The search is based on projecting the data onto wavelets of constant *quality factor*,

$$Q = \frac{f_c}{\sigma_f}, \tag{2}$$

for central frequency f_c and bandwidth σ_f. This projection was never tuned for the low frequencies required for detecting seismic noise, an example result is shown in Figure 4(a).

We have thus seen that the standard veto construction methods, that have proven effective for other systems, are less effective for seismic noise. This was the motivation for a different approach, resulting in the work described below.

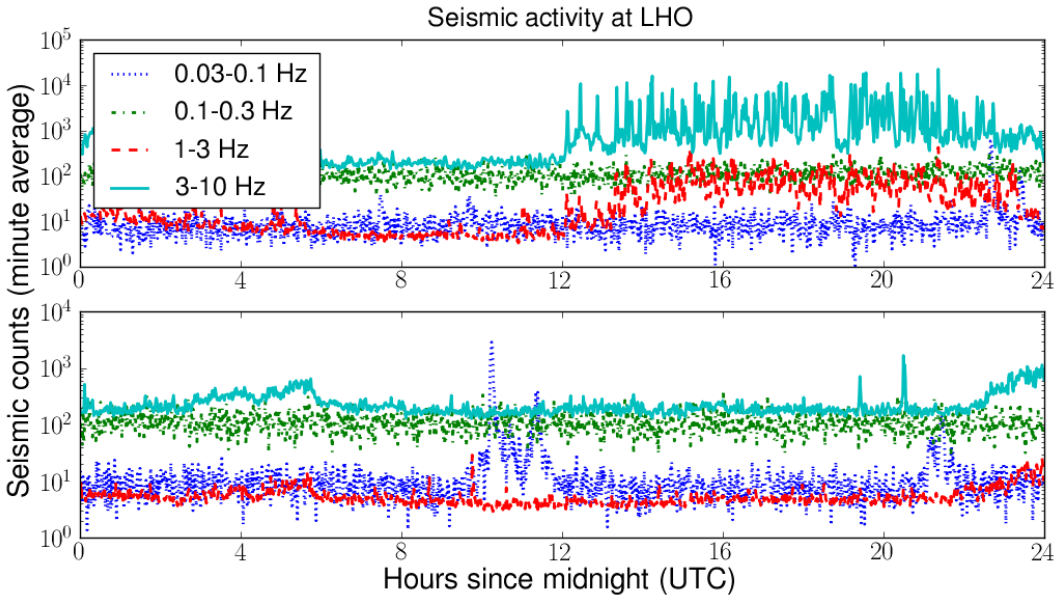

Figure 2. Examples of the difference between seismic noise incident at the LHO site during a full 24-hour span on a weekday (top) and a weekend (bottom). The difference in the higher frequency bands is the most pronounced.

4. Targeted veto methods

In order to produce effective vetoes, we have constructed a search algorithm to explicitly identify low-frequency seismic events, and construct lists of time segments to be excluded from GW searches. This method uses the *Ω pipeline* tuned specifically for low-frequency performance to generate lists of triggers highlighting seismic events, and the low-latency inspiral pipeline *Daily iHope* to generate lists of triggers from CBC template matched filtering. The two are combined by the *HierarchichalVeto* algorithm into lists of time segments during which seismic noise has polluted the GW analysis. In this section we will describe these three components, while in Section 5 we will present the results.

4.1. The Ω pipeline

The Ω-*Pipeline* is a burst detection algorithm developed within LIGO as a combination of the *Q Pipeline* [18] and *X-Pipeline* [19] and used, during S6, for low-latency detection of GW events to trigger electromagnetic followup [20]. The single-detector triggers were also used for DQ investigations.

The algorithm is based on the Q Transform [21] and projects detector data, $s(t)$, onto a bank of windowed complex exponentials of the following form:

$$S(\tau, f, Q) = \int_{-\infty}^{\infty} s(t) w(t-\tau, f, Q) \exp(-i2\pi f t)\, dt, \tag{3}$$

where w is a time-domain window centred on time τ, f is the central frequency, and Q is the quality factor. A three-dimensional bank is constructed in time-, frequency-, and

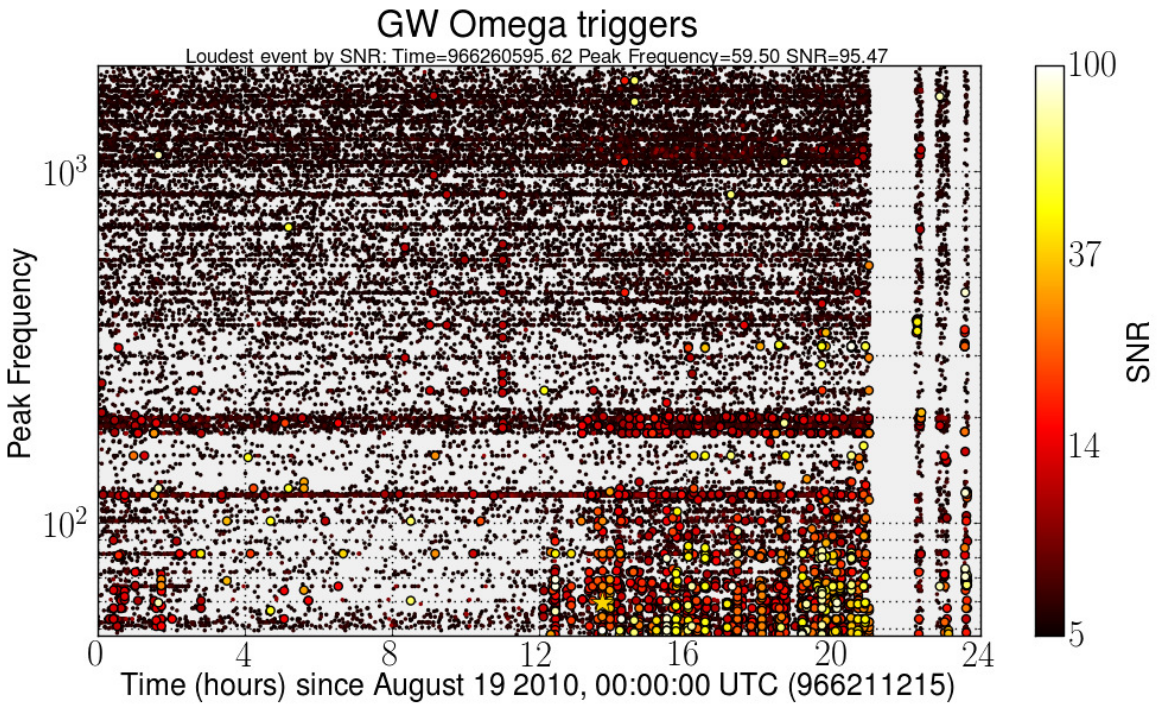

Figure 3. A 24-hour time-frequency map from the Ω-Pipeline of data from the GW readout. Hour 12 onwards are subject to increased noise at the lowest frequencies due to excess seismic noise from the working day. The columns of no events at the end of the period indicate that the detector was not operational.

Q-space such that the maximum loss in SNR as a result of the true signal lying between any two tiles does not exceed a given threshold. The *normalised energy*, Z, of a given tile is calculated as the ratio of its squared magnitude to the mean squared magnitude of other tiles in the same Q-frequency plane and is related to the SNR, ρ, as

$$\rho^2 = 2Z. \tag{4}$$

An example of the output of the Ω-Pipeline applied to low-latency gravitational wave readout data is show in Figure 3. A high density of low SNR (black) triggers is expected from Gaussian noise, but the higher SNR events (white) indicate increased noise at low frequencies, known to be correlated with higher seismic activity.

As a result of using both KW and Ω-Pipeline for DQ studies, direct comparisons were drawn on the performance of each, especially in frequency reconstruction at low frequency. It was found that, in their current implementations, the Ω-Pipeline gave much greater low frequency sensitivity, and better frequency resolution.

4.2. Paramaterisation of the Ω-Pipeline for seismic noise

The low-latency Ω-Pipeline analysis used a parameter set tuned for performance in the detection band, with a frequency range of 48–2048 Hz, and analyses done in 64 second blocks. As can be seen in Figure 3, the frequency range is such that the seismic band

is completely ignored, yet significant SNR is recorded up to $\sim$ 150 Hz that can be attributed to seismic noise.

In order to improve performance when applied to seismometer data, the parameter set applied to the algorithm was comprehensively tested, with the highest performance improvement seen from the following method of splitting the trigger generation into two bands. One set was trained on the anthropogenic band, above 2 Hz, and the other at very low frequency seismic activity, below 2 Hz. The following paragraphs detail the changes made to tune the Ω-Pipeline algorithm for each frequency band, describing three key parameters. The *sampling frequency* defines the time-resolution of the data to be filtered, and as such the maximum frequency reconstructed, and the *frequency range* gives the complete span of frequencies searched. The *block duration* defines the length of period used to estimate the power spectral density (PSD) of the detector.

4.2.1. The anthropogenic band, > 2 Hz

As described in Section 2 the seismic band extends upwards in frequency to around 30 Hz. Lowering the sampling frequency to 64 Hz‡ filtered out any high frequency seismic noise, allowing longer time-scale events to be triggered by the search. Also, as a result of dividing the parameter space into two parts, and lowering the sampling frequency for the anthropogenic band, the frequency range could be narrowed to run from 2 Hz up to the Nyquist frequency at 32 Hz. This allowed a great reduction in size of the bank of time-frequency-Q tiles, and reduced computational cost.

In the low-latency search for GWs using the Ω-Pipeline, blocks of 64 s were used to estimate detector sensitivity. In order to resolve events lasting up to a few seconds, the power spectrum was drawn from blocks of 4096 s. In this way a number of discrete seismic events above 2 Hz would not upset estimation of the power spectrum for the entire block.

4.2.2. The earthquake band, < 2 Hz

In this band, the sampling frequency could be reduced to 4 Hz, eliminating higher frequency disturbances and vastly reducing computational cost of the analysis. With this Nyquist frequency setting the upper bound on the frequency range, the lower bound was set to reflect the lowest sensible frequency given our knowledge of the seismic environment at both sites. As can be taken from Section 2 the lower limit was set to 0.01 Hz.

Finally, we have described that earthquake events typically dominate the seismic spectrum at either site (or even both sites simultaneously) for hours, requiring a background estimation at least an order of magnitude larger than this. As a result, blocks of 65536 seconds were used to capture the lower frequency noise within daily analyses.

‡ The Ω-Pipeline search for GWs downsamples the readout data to 4096 Hz, while the seismometers are only sampled at 256 Hz

Parameter	Untuned Value	Tuned value	
		< 2 Hz	> 2 Hz
Sample frequency	4096 Hz	4 Hz	64 Hz
Frequency range	$48 - 2048$ Hz	$0.01 - 2$ Hz	$2 - 32$ Hz
Block duration	64 s	65536 s	4096 s

Table 2. The parameter sets applied to the Ω-Pipeline search algorithm before and after tuning for low-frequency performance.

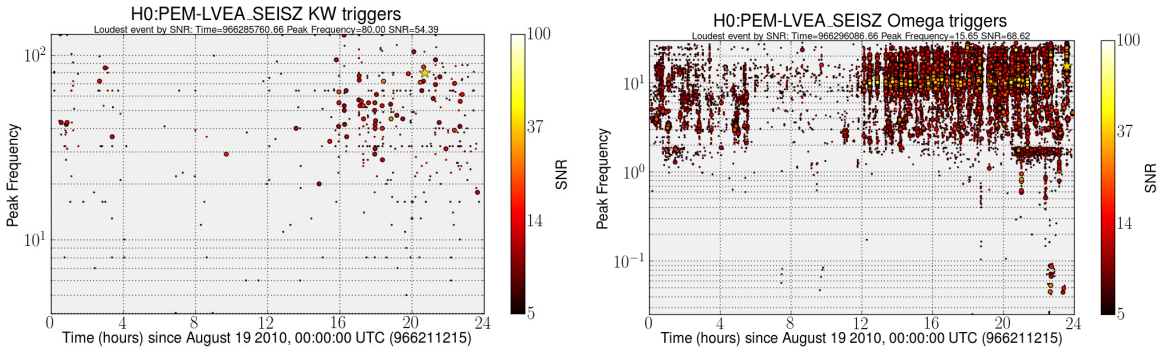

Figure 4. Examples of the untuned KW and tuned Ω-Pipeline algorithms applied to seismometer data. The left figure, 4(a), shows the lack of sensitivity in the untuned KW analysis, while that on the right, 4(b), has many orders of magnitude more events. Comparing to Figure 3 we can see loud (white) triggers around 10 Hz after 1200 UTC, but also triggers with SNR∼10 below 0.1 Hz.

As can be seen in Figure 4(b), the new parameter sets allows a huge increase in the number and significance of triggers produced by the Ω-Pipeline. The density of triggers can be greatly increased, especially around noisier times, with events recorded with frequencies as low as 0.03 Hz.

This method was applied to the four main seismometers at LHO: EX, EY, LVEA and VAULT§; and the three at LLO: EX, EY and LVEA (LLO has no VAULT seismometer).

4.3. Low-latency inspiral triggers – Daily iHope

The joint LIGO Virgo CBC group uses the *iHope* pipeline to search for GWs produced by binary coalescences. It is described more fully in [5, 22]. The seismic noise described above and shown in Figure 4(b) has been known to contribute to the noise background esimates in these searches, and so creating good vetoes specifically for CBC searches was a major goal of this work. Here we summarize the key points and discuss the changes

§ The EX and EY seismometers sit outside of the vacuum chambers containing the end test masses for the X- and Y-arms respectively, the Large Vacuum Equipment Area (LVEA) seismometer sits beside the chamber housing the GW readout photodetector, and the VAULT seismometer is in an underground chamber a small distance from the LVEA.

made for daily running of iHope in order to provide triggers to analyse alongside the Ω-Pipeline triggers from seismic data.

iHope is a templated, matched-filter search. In the low-mass ($M < 25\,\mathrm{M}_\odot$) region the template waveforms are restricted, stationary phase, frequency-domain waveforms of the form

$$\tilde{h}(f;M,\eta) = \frac{2GM_\odot}{c^2 r}\left(\frac{5M\eta}{96M_\odot}\right)^{\frac{1}{2}}\left(\frac{M}{\pi^2 M_\odot}\right)^{\frac{1}{3}} f^{-\frac{7}{6}}\left(\frac{GM_\odot}{c^3}\right)^{-\frac{1}{6}} e^{i\Psi(f;M,\eta)}, \tag{5}$$

where $M = m_1 + m_2$ is the total mass of the binary, $\eta = m_1 m_2/M^2$ is the symmetric mass ratio, and the phase evolution Ψ is known to 3.5 post-Newtonian (PN) order. A template bank is constructed by choosing a set of discrete points in the mass space at which to generate templates, and each of these is filtered individually. In the low-mass CBC search the total masses range from $2\,M_\odot - 25\,M_\odot$ and the set of parameters is chosen according to [23, 24] and the detector PSD. The bank is laid out such that the loss in SNR from a signal exactly matching the waveform family but with arbitrary parameters in the range is less than 3%.

To search for signals in the data we construct the matched filter. Denote the inner product of two waveforms h and s as

$$\langle s|h\rangle = 4\mathrm{Re}\int_{f_{\mathrm{low}}}^{f_{\mathrm{high}}} \frac{\tilde{s}(f)\tilde{h}^\star(f)}{S_n(f)}\,df, \tag{6}$$

where $S_n(f)$ is the PSD of the detector. The lower bound is taken to be 40 Hz, below which seismic noise dominates the detector output, as described in Sections 1 & 2. The upper bound is taken to be the frequency of the orbit of a point mass at the Schwarzschild innermost stable circular orbit (ISCO) above which the PN approximation becomes inaccurate.

Using the above inner product, the SNR is given by

$$\rho = \frac{\langle s|h\rangle}{\sqrt{\langle h|h\rangle}} \tag{7}$$

maximised over the phase. Triggers are selected by applying an SNR threshold of 5.5.

As mentioned above, the purpose of the daily iHope runs was not to search for GW, but rather to characterize the behavior of the detector. Therefore a number of simplifications were applied to the daily ihope template bank.

First, a static template bank was used for each interferometer, based on the layout at a quiet time in each instrument. Second, the low-mass end of the template bank is very dense‖, however, when looking for glitches such resolution is not needed, therefore

‖ Lower-mass systems have more cycles in the sensitive LIGO band, therefore there is ample information to distinguish between waveforms with close parameters

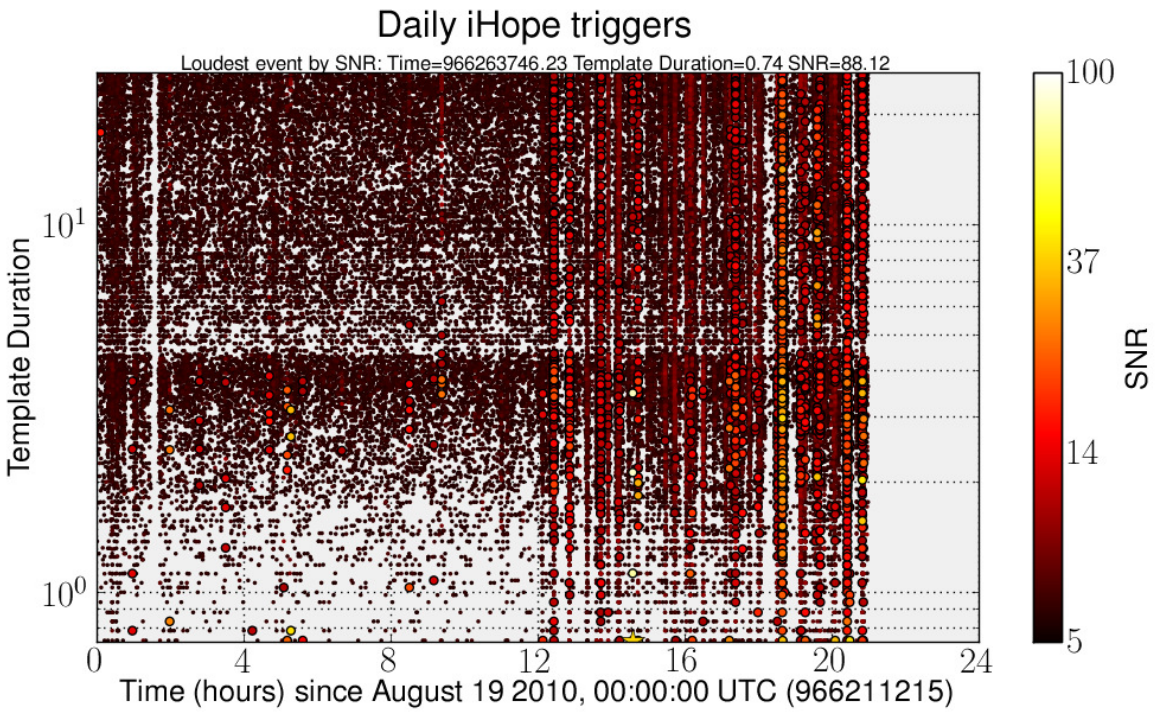

Figure 5. A 24-hour map of template duration versus time from the Daily iHope pipeline. CBC templates are characterised by a duration rather than a frequency, as they sweep through a range of frequencies. Comparing to Figures 3 and 2 shows the same excess of triggers after 1200 UTC, but here the excess noise at low frequency triggers across the entire template bank. No data is analyzed after 2100 because ihope requires at least 2048 contiguous seconds to estimate the PSD and all data after this time was in smaller segments.

the minimal match between templates in the region below a chirp mass ($\mathcal{M} = M\eta^{3/5}$) of $3.46\,M_\odot$ was set to 0.5. At higher masses, up to a total mass of $25\,M_\odot$ the minimal match was set to 0.95. In addition to ensuring coverage in a mass region which is naturally sparser this ensures that short glitches, characteristic of many glitch mechanisms, were flagged with large SNRs.

An example of the output of Daily iHope is shown in Figure 5.

4.4. Veto generation – HierarchichalVeto

The seismic triggers from the Ω-Pipeline, and the CBC triggers from Daily iHope were used to idenfity times of seismic noise using the statistical algorithm HierarchichalVeto (HVeto) [16]

The HVeto algorithm tests the statistical significance of time-coincidence between triggers from one channel, nominally the GW data channel, and those from auxiliary channels. The significance statistic is defined as

$$S_n(x) = -\log_{10}\left(\sum_{x}^{\infty} P_{poi}(\mu, x')\right), \tag{8}$$

where x' is the number of coincidences found between glitches in the auxiliary channel and the GW data channel, μ is the expected number of coincident events given the trigger rates in the two channels, x is the series of non-negative integers, and P_{poi} is the Poisson probability distribution function. The significance is calculated for all channels in a two-dimensional space of time-coincidence window, $T_{\rm win}$¶, and SNR threshold, ρ_c.

The most significant point on the $(T_{\rm win}, \rho_c)$ plane is chosen for each auxiliary channel, with the loudest channel by significance selected. Veto times are constructed by generating segments of width $T_{\rm win}$ around all triggers with SNR above ρ_c in that auxiliary channel. These segments are then removed from the analysis – allowing the next round to be 'won' by a (generally) different auxiliary channel containing less significant coincidences – and the procedure repeated until the significance of the loudest channel does not exceed a given stopping point. In this way, vetoes are generated hierarchichally, allowing for little redundancy between different channels.

Several modifications were made to this algorithm in order to test and run on the new seismic triggers. Significant testing was completed in order to construct a new $(T_{\rm win}, \rho_c)$ plane relevant for the long-duration events from the seismic data. Alongside this, as described in the caption to Figure 5, modifications were made to first read and understand the new Daily iHope triggers, and use the relevant new parameters.

5. Results

The method described above was that used in the construction of the LIGO seismic veto SeisVeto for S6. Here we present the results for a test sample of those data, spanning June 26 – August 6 2010. The results for LHO are shown in Table 2(a) and those for LLO in Table 2(b).

Each row in the table gives the statistics for the most significant channel in each frequency band, in addition to the cumulative results for the entire period$^+$.

For H1, each round contributes to a cumulative efficiency of 62.5 %, with a cumulative deadtime of 5.94 %. Simply put, almost two thirds of all triggers produced by the low-latency inspiral pipeline are occurring in a small amount of time, which is coincident with high seismic noise. This statistic alone outlines the problem caused by seismic noise.

It should not be surprising that the most significant channel for the two higher frequency bands should be the LVEA seismometer. This building is closer than any other to a major road, so experiences the highest magnitude of seismic noise from traffic and close anthropogenic noise, especially trucks serving the USDOE Hanford site, and

¶ Low-frequency events have a long duration whose maximum coupling time is not known. Also, many short-duration glitches in an auxiliary system can make take a certain time to couple into the GW output.

$^+$ The cumulative results include all rounds passing selection criteria in each band, not just the most significant channel.

(a) Results for H1.

Freq. Band (Hz)	Loudest Channel	Significance	Efficiency (%)	Deadtime (%)
0-1	EX	1455.21	3.26	0.15
1-3	EY	355.37	3.19	0.71
3-10	LVEA	12024.98	22.11	1.24
10-32	LVEA	41042.78	35.99	1.04
		Cumulative, all rounds	62.44	5.94

(b) Results for L1.

Freq. Band (Hz)	Loudest Channel	Significance	Efficiency (%)	Deadtime (%)
0-1	No channels passing selection criteria		0	0
1-3	LVEA	960.13	1.51	0.06
3-10	LVEA	420.55	0.88	0.06
10-32	EX	1601.22	2.29	0.07
		Cumulative, all rounds	6.95	0.60

Table 3. HVeto results from coincidence between new Ω-Pipeline triggers from seismometer data and low-latency inspiral triggers for the LIGO detectors operating in S6. Shown are the most significant (loudest) channels for each frequency band, and the cumulative statistics for the entire analysis.

also houses the majority of the interferometer control optics and subsystems, notably the GW readout photodetector.

For L1 we can see much lower statistical significance of the correlation between seismic noise and the readout signal. This can be attributed in part to the improvements from the HEPI feedforward system for the Livingston instrument, but also to the different nature of the seismic environment relative to LHO. This further highlights the problems faced on the Hanford site.

6. Future applications

One of the major improvements to be installed for the Advanced LIGO interferometers is the improved seismic isolation system [13]. By incorporating quadruple pendulum suspensions and additional active components this system will reduce the (linear) coupling of seismic noise into the interferometer readout by a significant factor (above 10 Hz), thus reducing the effect on the searches. The above veto generation method can be used to test the effectiveness of the new isolation, by accurately measuring the correlation between seismic noise (as measured by seismometers outside of the isolation) and the GW data.

Furthermore, as shown in section 5, targeting the trigger generator at a specific part of the parameter space relevant for a subsystem can greatly increase the power of statistical veto generation. This method is not only applicable to seismometer data, but can be used for any or all subsystems of the current, and next generation of LIGO

interferometers.

The seismic noise case was one for which the basic tuning of KW was especially poorly suited, and as such was one for which the targeted method was shown to give a massive improvement in efficiency and deadtime for vetoes constructed. In the aLIGO era, starting in the next few years, multiple instrumental and environmental systems can be characterised and targeted using this approach, with veto production for each system improved.

Maximising efficiency whilst minimising deadtime is a key factor in veto generation methods, and the above method has been seen to score well on both counts. This will lead to searches for GWs using cleaner data, with less wastage from long, inefficient vetoes.

7. Summary

We have highlighted the difficulties in generating effective seismic noise vetoes using the existing LIGO methods. The results presented have shown the benefits in using the targeted veto model we have introduced. The statistics alone demonstrate the magnitude of noise introduced by seismic activity, and so with a method to veto these times, search sensitivity will be improved.

Acknowledgements

The authors would like to thank Gabriela Gonzalez, Jessica McIver, Greg Mendell, Laura Nuttall, and all members of the Detector Characterization group, the Ω-Pipeline team, and the HierarchichalVeto team for discussions. DMM was supported by a studentship from the Science and Technology Facilities Council. SF was supported by the Royal Society. BH was supported by NSF grant PHY-0970074 and the UWM Research Growth Initiative. AL was supported by NSF grants PHY-0847611 (Syracuse) and PHY-0855589 (Penn. State). LP was supported by NSF grant PHY-0847611. JRS was supported by NSF grants PHY-0854812 (Syracuse) and PHY-0970147 (Fullerton). LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0107417.

[1] B. P. Abbott et al. LIGO: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, 72(7):076901, 2009.

[2] H. Grote et al. The status of GEO 600. *Classical and Quantum Gravity*, 22(10):S193, 2005.

[3] F. Acernese et al. Status of virgo. *Classical and Quantum Gravity*, 25(11):114045, 2008.

[4] J. R. Smith for the LIGO Scientific Collaboration. The path to the enhanced and advanced ligo gravitational-wave detectors. *Classical and Quantum Gravity*, 26, 2009.

[5] B. P. Abbott et al. Search for gravitational waves from low mass binary coalescences in the first year of LIGO's s5 data. *Phys. Rev. D*, 79(12):122001, Jun 2009.

[6] B. P. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev. D*, 80(10):102001, Nov 2009.

[7] L. Blackburn et al. The lsc glitch group: monitoring noise transients during the fifth LIGO science run. *Classical and Quantum Gravity*, 25(18):184004, 2008.

[8] N. Christensen and the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO s6 detector characterization studies. *Classical and Quantum Gravity*, 27(19):194010, 2010.

[9] J. Abadie et al. Characterization of the LIGO detectors during their sixth science run. *In preparation*, 2011.

[10] J. Slutsky et al. Methods for Reducing False Alarms in Searches for Compact Binary Coalescences in LIGO Data. *Class. Quant. Grav.*, 27:165023, 2010.

[11] E. J. Daw, J. A. Giaime, D. Lormand, M. Lubinski, and J. Zweizig. Long-term study of the seismic environment at LIGO. *Classical and Quantum Gravity*, 21(9):2255, 2004.

[12] S. Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New J. Phys.*, 11:123006, 2009.

[13] R. Abbott et al. Seismic isolation for advanced LIGO. *Classical and Quantum Gravity*, 19(7):1591, 2002.

[14] T. Isogai and the LIGO Scientific Collaboration and the Virgo Collaboration. Used percentage veto for LIGO and virgo binary inspiral searches. *Journal of Physics: Conference Series*, 243(1):012005, 2010.

[15] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21:1809–+, October 2004.

[16] J. R. Smith et al. A hierarchical method for vetoing noise transients in gravitational-wave searches. *In preparation*, 2011.

[17] Advanced LIGO. `http://www.ligo.caltech.edu/adv{LIGO}/`.

[18] S. K. Chattergi. *The search for gravitational wave bursts in data from the second LIGO science run.* PhD thesis, Massachusetts Instititute of Technology, 2005.

[19] P. J. Sutton et al. X-pipeline: an analysis package for autonomous gravitational-wave burst searches. *New Journal of Physics*, 12(5):053034, 2010.

[20] J. Adabie et al. Implementation and testing of the first prompt search for electromagnetic counterparts to gravitational wave transients. *In preparation*, 2011.

[21] J. C. Brown. Calculation of a constant q spectral transform. *The Journal of the Acoustical Society of America*, 89, 1991.

[22] B. P. Abbott et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run. *Phys. Rev. D*, 80:047101, 2009.

[23] B. J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, 1996.

[24] B. J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, 1999.